

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : G09G 5/00		A2	(11) International Publication Number: WO 98/49667
			(43) International Publication Date: 5 November 1998 (05.11.98)
(21) International Application Number: PCT/US98/06662		(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 1 April 1998 (01.04.98)			
(30) Priority Data: 08/848,569 28 April 1997 (28.04.97) US		Published Without international search report and to be republished upon receipt of that report.	
(71) Applicant: BOARD OF TRUSTEES, LELAND STANFORD, JR. UNIVERSITY [US/US]; Suite 350, 900 Welch Road, Palo Alto, CA 94304 (US).			
(72) Inventors: LEVOY, Marc; 18 Peter Coutts Circle, Stanford, CA 94305 (US). HANRAHAN, Pat; 40 Minoca Road, Portola Valley, CA 94028 (US).			
(74) Agents: MINSK, Alan, D. et al.; Coudert Brothers, Suite 3300, Four Embarcadero Center, San Francisco, CA 94111 (US).			

(54) Title: METHOD AND SYSTEM FOR LIGHT FIELD RENDERING

(57) Abstract

Described is a simple and robust method and system for generating new views from arbitrary camera positions without depth information or feature matching, simply by combining and resampling the available images. This technique interprets input images as two-dimensional slices of a four-dimensional function - the light field. This function completely characterizes the flow of light through unobstructed space in a static scene with fixed illumination. A sampled representation for light fields allows for both efficient creation and display of inward and outward looking views. Light fields may be created from large arrays of both rendered and digitized image. The latter are acquired with a video camera mounted on a computer-controlled gantry. Once a light field has been created, new views may be constructed in real time by extracting slices in appropriate directions. Also described is a compression system that is able to compress generated light fields by more than a factor of 100:1 with very little loss of fidelity. Issues of antialiasing during creation and resampling during slice extraction are also addressed.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

METHOD AND SYSTEM FOR LIGHT FIELD RENDERING

FIELD OF THE INVENTION

The present invention relates to the field of graphical imaging. More particularly, it relates to a system and method for generating new views of a three-dimensional scene from arbitrary positions without depth information or feature matching, by combining and resampling the available images.

BACKGROUND OF THE INVENTION

Traditionally, the input to a three-dimensional graphics system is a scene consisting of geometric primitives composed of different materials and a set of lights. Based on this input specification, the rendering system computes and outputs an image. Recently, a new approach to rendering has emerged which is known as image-based rendering. Image-based rendering systems generate different views of an environment from a set of pre-acquired imagery.

There are several advantages to the image-based rendering approach. First, the display algorithms for image-based rendering require modest computational resources and are thus suitable for real-time implementation on workstations and personal computers. Second, the cost of interactively viewing the scene is independent of the scene complexity. Even with the best algorithms and fastest hardware, it often takes minutes or hours using existing systems to generate images with the visual richness required for many applications. Third, the source of the pre-acquired images can be from a real or virtual environment (*i.e.* from digitized photographs or from rendered models), or from a combination of the two.

The forerunner to these techniques is the use of environment maps to capture the incoming light in a texture map. *See, e.g.,* J.F. Blinn & M.E. Newell, "Texture and Reflection in Computer Generated Images," *CACM*, Vol. 19, No. 10, Oct. 1976, pp. 542-47; N. Greene, "Environment Mapping and Other Applications of World Projections," *IEEE Computer Graphics and Applications*, Vol. 6, No. 11, Nov. 1986, pp. 21-29. An environment map records the incident light arriving from all directions at a point. The original use of environmental maps was to

- 2 -

efficiently approximate reflections of the environment on a surface. However, environment maps also may be used to quickly display any outward looking view of the environment from a fixed location but at a variable orientation. For example, this is the basis of the Apple QuickTime VR system. See S.E. Chen, "QuickTime VR -- An Image-Based Approach to Virtual Environment Navigation," Proc. SIGGRAPH '95 (Los Angeles, California, Aug. 6-11, 1995), *Computer Graphics Proceedings, Annual Conference Series, 1995, ACM SIGGRAPH*, pp. 29-38. In this system environment, maps are created at key locations in the scene. The user is able to navigate discretely from location to location, and while at each location continuously change the viewing direction.

The major limitation of rendering systems based on environment maps is that the viewpoint is fixed. One way to relax this fixed position constraint is to use view interpolation. See S.E. Chen & L. Williams, "View Interpolation for Image Synthesis," Proc. SIGGRAPH '93 (Anaheim, California, Aug. 1-6, 1993) *Computer Graphics Proceedings, Annual Conference Series, 1993, ACM SIGGRAPH*, pp. 279-88; N. Green and M. Kass, "Approximating Visibility with Environment Maps, Apple Technical Report no. 41, nov. 1994; H. Fuchs et al., "Virtual Space Teleconferencing Using a Sea of Cameras," *Proc. First International Conference on Medical Robotics and Computer Assisted Surgery*, 1994, pp. 161-67; L. McMillan & G. Bishop, "Head-Trackd Stereoscopic Display Using Image Warping," *Stereoscopic Displays and Virtual Reality Systems II*, Proc. SPIE, Vol. 240, S. Fisher, J. Merritt, B. Bolas eds. 1995, pp. 21-30; L. McMillan & G. Bishop, Plenoptic Modeling: An Image-Based Rendering System, Proc. SIGGRAPH '95 (Los Angeles, CA, Aug. 6-11, 1995) *Computer Graphics Proceedings, Annual Conference Series, 1995, ACM SIGGRAPH*, pp. 39-46 (hereinafter "McMillan, Plenoptic Modeling"); P.J. Narayana, "Virtualized Reality: Concepts and Early Results," *Proc. IEEE Workshop on the Representation of Visual Scenes*, IEEE, 1995. Most of these methods require a depth value for each pixel in the environment map, which is easily provided if the environment maps are synthetic

images. Given the depth value, it is possible to reproject points in the environment map from different vantage points to warp between multiple images. The key challenge in this warping approach is to "fill in the gaps" when previously occluded areas become visible.

5 Another approach to interpolating between acquired images is to find corresponding points in the two. See S. Laveau & O.D. Faugeras, "3-D Scene Representation as a Collection of Images and Fundamental Matrices," INRIA Technical Report No. 2205, 1994; McMillan, Plenoptic Modeling"; S. Seitz & C. Dyer, "Physically-Valid View Synthesis by Image Interpolation," *Proc. IEEE Workshop on the Representation of Visual Scenes*, IEEE, 1995. If the positions of
10 the cameras are known, this is equivalent to finding the depth values of the corresponding points. Automatically finding correspondence between pairs of images is the classic problem of stereo vision; unfortunately, although many algorithms exist, these algorithms are fairly fragile and may not always find the
15 correct correspondences.

 An abstract representation of light that is related to the present invention is epipolar volumes. See R. Bolles et al, "Epipolar-Plane Image Analysis: An Approach to Determining Structure from Motion," *International Journal of Computer Vision*, Vol. 1, No. 1, 1987, pp. 7-55. An epipolar volume is formed of
20 an array of images created by translating a camera in equal increments in a single direction. Such a representation has been used recently to perform view interpolation. Katayama et al. "Viewpoint-Dependent Stereoscopic Display Using Interpolation of Multiviewpoint Images," *Stereoscopic Displays and Virtual Reality Systems II*, Proc. SPIE, Vol. 2409, S. Fisher, J. Merrit, B. Bolas eds. 1995, pp. 11-
25 20.

 Another related representation is the horizontal-parallax-only holographic stereogram. S. Benton, "Survey of Holographic Stereograms," *Processing and Display of Three-Dimensional Data*, Proc. SPIE, Vol. 367, 1983. A holographic

5 stereogram is formed by exposing a piece of film to an array of images captured by a camera moving sideways. Halle has discussed how to set the camera aperture to properly acquire images for holographic stereograms. M. Halle, "Holographic Stereograms as Discrete Imaging Systems," *Practical Holography*, Proc. SPIE, Vol. 2176, Feb. 1994. Gavin Miller has also recognized the potential synergy between true three-dimensional display technologies and computer graphics algorithms. G. Miller, "Volumetric Hyper-Reality: A Computer Graphics Holy Grail for the 21st Century?" *Proc. Graphics Interface '95*, W. Davis and P. Prusinkiewicz eds., Canadian Information Processing Society, 1995, pp. 56-64.

- 5 -

SUMMARY OF THE INVENTION

The primary object of this invention is to provide a technique for displaying three dimensional scenes from previously acquired imagery. A term commonly applied to such techniques is image-based rendering. It is a further object of the invention to provide a new image-based rendering technique which is robust and allows much more freedom in the range of possible views, particular in comparison with the above described techniques.

In fulfillment of these objects and others, our technique utilizes a representation of the "light field," a term used herein to denote the radiance as a function of position and direction in regions of space free of occluders ("free space"). In free space, the light field is a four-dimensional, rather than a five-dimensional function. An image is a two-dimensional slice of the four-dimensional light field. Creating a light field from a set of images corresponds to inserting each two-dimensional slice into the four-dimensional light field representation. Similarly, generating new views corresponds to extracting and resampling a slice.

It will be appreciated by those skilled in the art that generating a new image from a light field is quite different than previous view interpolation approaches. First, the new image is generally formed from many different pieces of the original input images, and need not look like any of them. Second, no model information, such as depth values or image correspondences, is needed to extract the image values. Third, image generation involves only resampling, a simple linear process. Fourth, this impementation may require a large amount of data, mandating some sort of compression.

As described in greater detail below, when the invention is implemented to display three-dimensional scenes on a graphical workstation, it may be necessary to address several issues. First, choice of proper parameterization and representation of the light field is required as is the choice of sampling pattern for the field.

- 6 -

Second, one must determine how to generate or acquire the light field. Third, one must address problems related to the fast generation of different views. This requires that a slice representing rays through a point be easily extracted, and that the slice be properly resampled to avoid artifacts in the final image. Fourth, this implementation may require a large amount of data. A preferred technique for addressing these issues is described below.

Additional objects, advantages and novel features of the invention will become more apparent to those skilled in the art upon examination of the following description or may be learned by practice of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and form a part of the specification, illustrate the present invention and, together with the description, serve to explain the principles of the invention. In the drawings:

Fig. 1 illustrates a light slab representation.

Figs. 2A and 2B illustrate the representation of an oriented line in Cartesian space and in line space.

Figs. 3A to 3D illustrate the use of multiple light slabs to represent various views of an object.

Figs. 4A to 4D illustrate the density of samples in line space for different arrangements of light slabs.

Fig. 5 illustrates the viewing geometry used to create light slabs from an array of perspective images.

Figs. 6A and 6B show two visualizations of a light field.

Fig. 7 illustrates a technique for prefiltering a light field.

Fig. 8 illustrates a prefiltering technique using an aperture.

- 7 -

Fig. 9 is a perspective view of a camera gantry used to obtain digitized images in accordance with an embodiment of the invention.

Fig. 10 is a perspective view of an object and lighting support system for digitizing images in accordance with an embodiment of the invention.

5 Fig. 11 a representation of a two-stage compression pipeline.

Fig. 12 illustrates a technique for resampling a light slab during display in accordance with the invention.

Fig. 13 illustrates of the effects of interpolation during slice extraction.

Fig. 14 shows example images from four light fields.

10 Fig. 15 shows images extracted from compressed light fields.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Using the drawings, the preferred embodiment of the present invention will now be explained. This discussion first describes the representation of a light field. It then describes the creation of light fields from rendered images and from
15 digitized images. A preferred compression technique for use in connection with the invention is then described. A technique for displaying an image from a light slab is then described. Finally, illustrative images are described which were constructed and extracted in accordance with the invention.

A. Representation

20 As used herein, light field is defined as the radiance at a point in a given direction. This definition is equivalent to the plenoptic function introduced by Adelson and Bergen. See E.H. Adelson & J.R. Bergen, "The Plenoptic Function and the Elements of Early Vision," *Computation Models of Visual Processing*, M.

Landy and J.A. Movshon, eds., MIT Press, Cambridge, 1991. The phrase light field was coined by A. Gershun in his paper, "The Light Field," Moscow, 1936, *Journal of Mathematics and Physics*, P. Moon & G. Timoshenko, trans., Vol. XVIII, MIT, 1939, pp. 51-151.

5 McMillan and Bishop discuss the representation of five-dimensional light fields as a set of panoramic images at different three-dimensional locations. See McMillan, *Plenoptic Modeling*, *supra*. However, the five-dimensional representation may be reduced to four-dimensions in free space. This is a consequence of the fact that the radiance does not change along a line unless
10 blocked. Four-dimensional light fields may be interpreted as functions on the space of oriented lines. The redundancy of the five-dimensional representation is undesirable for two reasons: first, redundancy increases the size of the total dataset, and second, redundancy complicates the reconstruction of the radiance function from its samples. This reduction in dimension has been used to simplify
15 the representation of radiation emitted by luminaires. See R. Levin, "Photometric Characteristics of Light Controlling Apparatus," *Illuminating Engineering*, Vol. 66, No. 4, 1971, pp. 205-15; I. Ashdown, "Near-Field Photometry: A New Approach," *Journal of the Illuminating Engineering Society*, Vol. 22, No. 1, Winter 1993, pp. 163-80. Accordingly, the following discussion is directed to four-dimensional light
20 fields.

 Although restricting the validity of the representation to free space may seem like a limitation, there are two common situations where this assumption is useful. First, most geometric models are bounded. In this case, free space is the region outside the convex hull of the object, and hence all views of an object from
25 outside its convex hull may be generated from a four-dimensional light field. Second, if one moves through an architectural model or an outdoor scene, one usually moves through a region of free space, and therefore, any view from inside this region of objects outside the region may be generated.

- 9 -

A representation of a four-dimensional light field must be selected to parameterize the space of oriented lines based on several considerations, including: (1) calculation efficiency; (2) control over the set of lines; and (3) uniform sampling. The first consideration requires that the computation of the position of a line from its parameters be fast. More importantly, for the purposes of calculating new views, it should be easy to compute the line parameters given the viewing transformation and a pixel location. The second consideration is based on the realization that while the space of all lines is infinite, only a finite subset of line space is ever needed. For example, in the case of viewing an object we need only lines intersecting the convex hull of the object. Thus, there is an intuitive connection between the actual lines in three-dimensional space and line parameters. The third consideration reflects that given equally spaced samples in line parameter space, the pattern of lines in three-dimensional space should also be uniform. In this sense, a uniform sampling pattern is one where the *number of lines* in intervals between samples is constant everywhere.

The present invention utilizes a representation which addresses these considerations. According to this representation, lines are parameterized by their intersections with two planes in arbitrary position. As illustrated in Fig. 1, a convention is adopted wherein the coordinate system on a first plane 2 is (u, v) and on a second plane 4 is (s, t) . An oriented line 6 is defined by connecting a point 8 on the $u-v$ plane to a point 10 on the $s-t$ plane. In the preferred embodiment u , v , s , and t are restricted to lie between 0 and 1, and thus points on each plane are restricted to lie within a convex quadrilateral. This representation is referred to herein as a "light slab." Intuitively, a light slab represents the beam of light entering one quadrilateral and exiting another quadrilateral.

It is an advantageous feature of this representation that one of the planes may be placed at infinity. This is convenient since it permits lines to be parameterized by two points, or by a point and a direction. The latter proves useful

- 10 -

for constructing light fields from orthographic images. Furthermore, if all calculations are performed using homogeneous coordinates, the two cases may be handled at no additional cost in computation.

Another significant advantage of this representation is the efficiency of geometric calculations. Mapping from (u,v) to points on the plane is a projective map and involves only linear algebra (multiplying by a 3×3 matrix). More importantly, as will be discussed below, the inverse mapping from an image (x,y) to (u,v,s,t) is also a projective map. Methods using spherical or cylindrical coordinates require substantially more computation.

Figs. 2A and 2B use line space to illustrate some of the major concepts of our representation. In these figures, lines are shown in two dimensions to simplify visualization. It will be appreciated that the extension to three dimensions is straightforward.

In Fig. 2A, a line 20 in Cartesian space is shown. As shown in Fig. 2B, the line 20 may be represented in line space by a point. Thus, each oriented line may be represented in line space by a corresponding point. It will be appreciated that many properties of light fields are easier to understand in line space. In line space, each oriented line is represented by a point and each set of lines by a region. In particular, the set of lines represented by a light slab and the set of lines intersecting the convex hull of an object are both regions in line space. All views of an object could be generated from one light slab if its set of lines completely surrounded the set of lines intersecting the convex hull of the object. Unfortunately, this is not possible. Therefore, it takes multiple light slabs to represent all possible views of an object. The system therefore tiles line space with a collection of light slabs.

Figs. 3A to 3D illustrate the use of multiple light slabs to represent various views of an object. Fig. 3A shows a single light slab 30. Light rays 32 connect points on two defining lines 34 and 36. Fig. 3B shows the corresponding line

space visualizations of the light rays shown in Fig. 3A. Fig. 3C shows an arrangement 35 of four rotated copies 30a, 30b, 30c, 30d of the light slab 30 shown in Fig. 3A. Fig. 3D shows the corresponding line space visualizations for each light ray of these rotated copies.

5 For any set of lines in Cartesian space, the envelope formed by the corresponding points in line space indicates the coverage of position and direction. Ideally, the coverage should be complete in θ and as wide as possible in r . As Figs. 3A to 3D show, the single slab 30 in Fig. 3A does not provide full coverage in θ , but the four-slab arrangement 35 in Fig. 3C does. The representation of Fig. 3C, however, is narrow in r . Such an arrangement is suitable for inward-looking views of a small object placed at the origin. Illustratively, it was used to generate the light field of the lion object described below in reference to Fig. 14D.

10 An important issue related to the parameterization is the sampling pattern. Assuming that all views are equally likely to be generated, then any line is equally likely to be needed. Thus all regions of line space should have an equal density of samples. Figs. 4A to 4D show the density of samples in line space for different arrangements of slabs.

15 Fig. 4A shows a light slab defined by two lines 42 and 44 at right angles. Fig. 4C shows a light slab where one defining line is at infinity and the other line 46 lies along the x -axis. This arrangement generates rays passing through the defining line 46 with an angle between -45° and $+45^\circ$. Fig. 4B shows the line space visualization corresponding to the light slab of Fig. 4A. Fig. 4D shows the line space visualization corresponding to the light slab of Fig. 4C.

20 The use of (r, θ) to parameterize line space has the property that equal areas in line space correspond to equally dense sampling of position and orientation in Cartesian space; ideally the density of points in line space should be uniform. As Figs. 4A to 4D show, the singularity at the corner 48 in Fig. 4A leads to a highly

- 12 -

nonuniform, and, therefore, inefficient sampling pattern, indicated by dark areas 43, 47 in Fig. 4B at angles of 0 and $-\pi/2$. The light slab shown in Fig. 4C generates a more uniform set of lines. Although Fig. 4C does not provide full coverage of θ , four rotated copies do. Such an arrangement is suitable for outward-looking views by an observer standing near the origin. For example, this arrangement may be used to generate the hallway light field in Fig. 14C.

The representations shown in Figs. 3A to 3D and Figs. 4A to 4D demonstrate that no slab arrangement is perfect. For example, as shown in Fig. 4A, slab arrangements with a singularity, such as two polygons joined at a corner, are bad and should be avoided. In contrast, as shown in Fig. 3A, slabs formed from parallel planes generate fairly uniform patterns. In addition, arrangements like the one shown in Fig. 4C are better than those with two finite plane (such as those shown in Fig. 3A). Finally, because of symmetry, the spacing of samples in u - v should be roughly the same as s - t . If the observer is likely to stand near the u - v plane, then it may be acceptable to sample u - v less frequently than s - t .

B. Creation of Light Fields

The creation of both virtual light fields (from rendered images) and real light fields (from digitized images) are now discussed in reference to Figs. 5 to 10. One method to create a light field is to choose a four-dimensional sampling pattern, and for each line sample, find the radiance. This is easily done directly for virtual environments by a ray tracer. In a real environment the preferred approach to generate light fields is to assemble a collection of images of the environment. Alternatively, a spot radiometer may be used, but this approach may be tedious.

1. Creation of Light Fields from Rendered Images

For a virtual environment, a light slab is easily generated simply by rendering a two-dimensional array of images. Each image represents a slice of the

- 13 -

four-dimensional light slab at a fixed $u-v$ value and is formed by placing the center of projection of the virtual camera at the location of the sample on the $u-v$ plane. The only issue is that the $x-y$ samples of each image must correspond exactly with the $s-t$ samples. As described below, this is easily done by performing a sheared perspective projection similar to that used to generate a stereo pair of images.

Fig. 5 illustrates the viewing geometry to create a light slab from an array of perspective images. Each perspective is arranged at positions 52a-52d along the camera plane 54, which corresponds to the $u-v$ plane utilizing the convention adopted herein. The focal plane 56 corresponds to the $s-t$ plane.

Figs. 6A and 6B show the resulting four-dimensional light field utilizing the viewing geometry shown in Fig. 5. This light field can be visualized either as a $u-v$ array of $s-t$ images (Fig. 6A) or as an $s-t$ array of $u-v$ images (Fig. 6B). In Fig. 6A, each image 60 in an array 62 represents the rays arriving at one point 64 on the $u-v$ plane from all points on the $s-t$ plane. In Fig. 6B, each image 65 in array 66 represents the rays leaving one point 68 on the $s-t$ plane bound for all points on the $u-v$ plane. The images in Fig. 6A are off-axis (*i.e.* sheared) perspective views of the scene, while the images in Fig. 6B look like reflectance maps. The latter occurs because the object has been placed astride the focal plane, making sets of rays leaving points on the focal plane similar in character to sets of rays leaving points on the object.

Alternatively, other viewing geometries are useful. For example, a light slab may be formed from a two-dimensional array of orthographic views. This can be modeled by placing the $u-v$ plane at infinity, as shown in Fig. 4A. In this case, each $u-v$ sample corresponds to the direction of a parallel projection. Again, the only issue is to align the $x-y$ and $s-t$ samples of the image with the $s-t$ quadrilateral. Another useful geometry consists of a two-dimensional array of outward looking (non-sheared) perspective views with fixed field of view. In this case, each image

- 14 -

is a slice of the light slab with the s - t plane at infinity. The fact that all these cases are equally easy to handle with light slabs attests to the elegance of projective geometry. Light fields using each arrangement are described below and illustrated in Fig. 14.

5 As with any sampling process, sampling a light field may lead to aliasing since typical light fields contain high frequencies. Fortunately, the effects of aliasing may be alleviated by filtering before sampling. In the case of a light field, a four-dimensional filter in the space of lines is employed. Such a filter is shown in Fig. 7 where a ray filter 76 is produced by the combination of a pixel filter 72 and
10 an aperture filter 74.

One possible embodiment employs a box filter. In this case, a weighted average of the radiances on all lines connecting sample squares in the u - v and s - t planes is computed. Light fields are then created from images generated by a camera placed on the u - v plane which is focused on the s - t plane. The filtering
15 process corresponds to integrating both over a pixel corresponding to an s - t sample, and an aperture equal in size to a u - v sample. This technique is described in reference to Fig. 8.

Fig. 8 illustrates a prefiltering technique using an aperture. As shown, a camera focuses on the s - t plane 86. A hypothetical film plane 85 is drawn behind
20 the aperture 82 which is disposed in the u - v plane 84. Ignoring the aperture momentarily (for example, in the case of a pinhole camera that precisely images the s - t plane onto the film plane 85), integration over a pixel on the film plane is then equivalent to integrating over an s - t region bounded by the pixel. Next, considering
25 a fixed point on the film plane while using finite sized aperture (all rays from a point on the film through the aperture are focused on a single point on the focal plane), integration over the aperture 82 corresponds to integrating all rays through the u - v region bounded by the aperture. Therefore, by simultaneously integrating

- 15 -

over both the pixel and the aperture, the proper four-dimensional integral is computed. The theory behind this filtering process has been discussed in the context of holographic stereograms by Halle. See M. Halle, "Holographic Stereograms as Discrete Imaging Systems," *Practical Holography*, Proc. SPIE, Vol. 2176, Feb. 1994.

Although prefiltering has the desired effect of antialiasing the light field, it has what at first seems like an undesirable side effect -- introducing blurriness due to depth of field. However, this blurriness is precisely correct for the situation described herein. In particular, when creating a pair of images from two adjacent camera locations on the u - v plane, a given object point will project to different locations in these two images. The distance between the two projected locations is called the stereo disparity. Extending this idea to multiple camera locations produces a sequence of images in which the object appears to jump by a distance equal to the disparity. This jumping is aliasing. Taking an image with a finite aperture causes points out of focus to be blurred on the film plane by a circle of confusion. Setting the diameter of the aperture to the spacing between camera locations causes the circle of confusion for each object point to be equal in size to its stereo disparity. This replaces the jumping with a sequence of blurred images. Thus, aliasing is removed by employing finite depth of field.

The necessity for prefiltering can also be understood in line space. Since samples of the light field correspond to points in line space, having a finite depth of field with an aperture equal in size to the u - v sample spacing insures that each sample adequately covers the interval between these line space points. Too small or too large an aperture yields gaps or overlaps in line space coverage, resulting in views that are either aliased or excessively blurry, respectively.

- 16 -

2. Digitized Images

Digitizing the imagery required to build a light field of a physical scene is a formidable engineering problem. The number of images required is large (hundreds or thousands), so the process preferably is automated or at least computer-assisted. Moreover, the lighting must be controlled to insure a static light field, yet flexible enough to properly illuminate the scene, all the while staying clear of the camera to avoid unwanted shadows. Finally, real optical systems impose constraints on angle of view, focal distance, depth of field, and aperture, all of which must be managed. Similar issues have been faced in the construction of devices for performing near-field photometric measurements of luminaries. I. Ashdown, "Near-Field Photometry: A New Approach," *Journal of the Illuminating Engineering Society*, Vol. 22, No. 1, Winter 1993, pp. 163-80. The following discussion describes the issues related to digitizing images and our preferred approach (discussed in reference to Figs. 9 and 10) for addressing them. Based on this discussion, various other approaches will be apparent to those skilled in the art.

a. Inward versus outward looking

Imagery may be based on either a flyaround of a small object or a flythrough of a large-scale scene. In the embodiment described in this section, the former approach was utilized. Thus, the generated imagery is inward looking. For example, the lion light field of Fig. 14D may be generated with the system shown in Fig. 9.

b. Human versus computer-controlled

An inexpensive approach to digitizing light fields is to move a handheld camera through the scene, populating the field from the resulting images.

However, this approach necessitates estimating camera pose at each frame and interpolating the light field from scattered data -- two challenging problems. As a simplification, the preferred embodiment described herein employs a computer-

- 17 -

controlled camera gantry and digitized images on a regular grid. Such a system is shown in Figs. 9 and 10.

c. Spherical versus planar camera motion

Many possible camera motions lend themselves to creating the images
5 needed to construct a light field. As examples, we compare two such motions:
motion along the surface of a sphere, or motion in a plane. These lead naturally to
implementations involving a spherical gantry or a planar gantry, respectively. A
spherical gantry has three advantages over a planar gantry: (a) it is easier to cover
the entire range of viewing directions; (b) the sampling rate in direction space is
10 more uniform; and (c) the distance between the camera and the object is fixed,
providing sharper focus throughout the range of camera motion. On the other
hand, a planar gantry has two advantages over a spherical gantry: (a) it is easier to
build (the entire structure can be assembled from linear motion stages); and (b) it is
closer to the light slab representation described above.

15 In view of these considerations, the embodiment shown in Fig. 9 employs a
planar gantry to obtain images of an object 96. While various components may be
utilized, the illustrative embodiment shown includes a modified Cyberware MS
motion platform with additional stepping motors which provide four degrees of
freedom (horizontal translation, vertical translation, pan and tilt). A Panasonic
20 WV-F300 3-CCD video camera 92 is used with a Canon f/1.7 10-120 mm zoom
lens. In use, the lens may be locked at its widest setting (10mm) and mounted so
that the pitch and yaw axes pass through the center of projection. While digitizing,
the camera 92 is kept pointed at the center of the focal plane. Calibrations and
alignments are verified with the aid of a Faro digitizing arm 94 which is accurate to
25 0.3 mm.

An alternative implementation, in which the camera is moved along the
surface of a sphere, might be a gantry design consisting of two concentric

- 18 -

hemicycles similar to a gyroscope mounting. Such a gantry has been constructed by Apple Computer to acquire imagery for Quick-Time VR flyarounds. See S.E. Chen, "QuickTime VR -- An Image-Based Approach to Virtual Environment Navigation," Proc. SIGGRAPH '95 (Los Angeles, California, Aug. 6-11, 1995), *Computer Graphics Proceedings, Annual Conference Series, 1995, ACM SIGGRAPH* pp. 29-38. However, the lighting in the Apple Computer system is attached to the moving camera, rendering it unsuitable for acquiring static light fields.

d. Field of view

Preferably, a light field is built that allows 360 degrees of azimuthal viewing. To accomplish this using a planar gantry involves acquiring four slabs each providing 90 degrees.

The example shown in Fig. 9 achieves this objective with a camera equipped with pan and tilt motors, thereby enabling the use of a narrow-angle lens. The use of a rotating camera means that, in order to transfer the acquired image to the light slab representation it must be reprojected to lie on a common plane. This reprojection is equivalent to keystone correction in architectural photography.

Alternatively, one may use a camera that translates but does not pan or tilt by employing wide-angle optics. However, this solution has two disadvantages. First, wide-angle lenses exhibit significant distortion which must be corrected after acquisition. Second, this solution trades off angle of view against sensor resolution.

Yet another alternative is to employ a view camera in which the sensor and optical system translate in parallel planes, the former moving faster than the latter. For example, horizontal parallax holographic stereograms are constructed using such a camera. See M. Halle, "Holographic Stereograms as Discrete Imaging

- 19 -

Systems," *Practical Holography*, Proc. SPIE, vol. 2176, Feb. 1994. However, incorporating such a solution into a gantry that moves both horizontally and vertically may prove difficult.

e. Standoff distance

5 A disadvantage of planar gantries is that the distance from the camera to the object changes as the camera translates across the plane, making it difficult to keep the object in focus. The view camera 92 described in reference to Fig. 9 does not suffer from this problem because the ratio of object distance to image distance stays constant as the camera translates. For the rotating camera 92 described in reference
10 to Fig. 9, servo-controlled focusing is an option, but changing the focus of a camera shifts its center of projection and changes the image magnification, complicating acquisition. This problem may be minimized by using strong lighting and a small aperture to maximize depth of field.

f. Sensor rotation

15 Each sample in a light slab should represent the integral over a pixel, and these pixels should lie on a common focal plane, as discussed above. A view camera satisfies this constraint because its sensor translates in a plane. The use of a rotating camera means that the focal plane also rotates. However, if the images are resampled carefully during reprojection, the presence of a rotated focal plane will
20 introduce no additional error into the light field. In the preferred embodiment, no artifacts were observed due to such a resampling process.

g. Aperture size

25 Each sample in a light slab should also represent the integral over an aperture equal in size to a u - v sample. In the system shown in Fig. 9, the use of a small aperture produces a light field with little or no u - v antialiasing. Even fully open, the apertures of many commercial video cameras are small. The required

- 20 -

antialiasing may be approximated by averaging together some number of adjacent views, thereby creating a synthetic aperture. However, this technique requires a very dense spacing of views, which in turn requires rapid acquisition.

h. Object support

5 In order to acquire a 360-degree light field in four 90-degree segments using a planar gantry, either the gantry or the object must be rotated to each of four orientations spaced 90 degrees apart. In view of the size of the gantry illustrated in Fig. 9, the latter technique was adopted in this example. This technique is illustrated in reference to Fig. 10.

10 Fig. 10 illustrates an object and lighting support. In this example, objects are mounted on a platform 108 attached to a tripod 102, such as a Bogen fluid-head tripod which can be rotated to four orientations spaced 90 degrees apart.

i. Lighting

15 As illustrated in Fig. 10, a rotating lighting system 104, 106 may be used to satisfy the requirement for fixed illumination. Illumination is provided by lights 104 attached to a ceiling mounted rotating hub 106 that is aligned with the rotational axis of the tripod 102. Illustratively, the lights 104 may comprise two 600W Lowell Omni spotlights. Preferably, a stationary diffuser panel (not shown) is hung between the spotlights and the gantry and the entire apparatus enclosed in black
20 velvet to eliminate any stray light. Alternatively, a lighting system which exhibits fourfold symmetry could also be used.

3. Methodology for Acquiring a Light Field

Utilizing the system described above in reference to Figs. 9 and 10, it is possible to acquire a light field according to the following procedure. This
25 procedure comprises the following steps. First, for each of four orientations, the

- 21 -

camera is translated through a regular grid of camera positions. Next, the camera is panned and tilted at each position to point at the center of the object, along the axis of rotation of the tripod. The image is then acquired, and using standard texture mapping algorithms known in the art, the image is reprojected to lie on a common plane as described above.

Table II, set forth below, provides an illustrative set of acquisition parameters. In that example, the distance between the camera positions (3.125 cm) exceeds the diameter of the aperture (1.25 mm). Accordingly, this example could benefit from denser spacing and a synthetic aperture as noted above.

C. Data Compression

Light field arrays are large, and to make creation, transmission, and display of light fields practical, they preferably should be compressed. Several compression techniques are available. However, we have discovered that the following unique characteristics of light fields are significant in determining a proper compression technique for use in connection with the invention.

First, data coherence is a significant characteristic of light fields with respect to compression. Light fields exhibit coherence in all four dimensions. For example, the smooth regions in Fig. 6A indicate that this light field contains coherence in s and t , and the smooth regions in Fig. 6B indicated that the light field contains coherence in u and v . The former corresponds to our usual notion of interpixel coherence in a perspective view. The latter can be interpreted either as the interframe coherence one expects in a motion sequence or as the smoothness one expects in the bidirectional reflectance distribution function (BRDF) for a diffuse or moderately specular surface. Occlusions introduce discontinuities in both cases.

- 22 -

Second, most compression techniques place some constraint on random access to data. For example, variable bitrate coders may require scanlines, tiles or frames to be decoded at once. More particularly, such coders include variable bitrate vector quantization and the Huffman or arithmetic coders used in JPEG or MPEG. Predictive coding schemes further complicate access because pixels depend on previously decoded pixels, scanlines, or frames. This poses a problem for light fields since the set of samples referenced when extracting an image from a light field are dispersed in memory. As the observer moves, the access patterns change in complex ways. Therefore, a compression technique that supports low-cost random access to individual samples is preferred.

Third, asymmetry is a characteristic which relates to a proper compression technique. Applications of compression can be classified as symmetric or asymmetric depending on the relative time spent encoding versus decoding. In a preferred embodiment of the invention, it is assumed that light fields are assembled and compressed ahead of time, making this an asymmetric application.

Fourth, computational expense is a factor in choice of a compression technique. Preferably, a compression scheme can be decoded without hardware assistance. Although software decoders have been demonstrated for standards like JPEG and MPEG, these implementations consume the full power of a modern microprocessor. In addition to decompression, the display method has additional work to perform as described more particularly below. Therefore, it is preferred that a compression technique be utilized which can be decoded quickly.

In view of these considerations, the preferred embodiment utilizes a two-stage pipeline compression scheme consisting of fixed rate vector quantization followed by entropy coding, such as Lempel-Ziv coding. This technique is illustrated in reference to Fig. 11.

- 23 -

As represented in Fig. 11, a light field 110 is partitioned into tiles, which are encoded using vector quantization (VQ) 112 to form an array of codebook indices. In this example, the codebook 114 and the array of indices 116 are further compressed using Lempel-Ziv (LZ) coding. An output bit stream 118 is then
5 obtained.

In this example, decompression occurs in two stages: entropy decoding as the file is loaded into memory, and dequantization on demand during interactive viewing. Typical file sizes for this example are shown in Fig. 11 beside each stage.

1. Vector Quantization

10 As noted above, the first stage of the compression pipeline shown in Fig. 11 is vector quantization (for a more complete description of this technique, see A. Gersho et al., *Vector Quantization and Signal Compression*, Kluwer Academic Publishers 1992). Vector quantization is a lossy compression technique wherein a vector of samples is quantized to one of a number of predetermined reproduction
15 vectors. A reproduction vector is called a codeword, and the set of codewords available to encode a source is called the codebook. Codebooks are constructed during a training phase in which the quantizer is asked to find a set of codewords that best approximates a set of sample vectors, called the training set. The quality of a codeword is typically characterized using mean-squared error (MSE) (*i.e.* the
20 sum over all samples in the vector of the squared difference between the source sample and the codeword sample). Once a codebook has been constructed, encoding consists of partitioning the source into vectors and finding for each vector the closest approximating codeword from the codebook. Decoding consists of looking up indices in the codebook and outputting the codewords found there -- a
25 very fast operation. Indeed, decoding speed is one of the primary advantages of vector quantization.

- 24 -

In the preferred embodiment, two-dimensional or four-dimensional tiles of the light field are used, yielding 12-dimensional or 48-dimensional vectors, respectively. The former takes advantage of coherence in s and t only, while the latter takes advantage of coherence in all four dimensions. To maximize image quality, a representative subset of each light field is trained to be compressed, and then the resulting codebook is transmitted along with the codeword index array. Since light fields are large, even after compression, the additional overhead of transmitting a codebook is small, typically less than 20%. By training on a subset rather than the entire light field, the expense of training is reduced.

The output of vector quantization is a sequence of fixed-rate codebook indices. Each index is $\log N$ bits where N is the number of codewords in the codebook, so the compression rate of the quantizer is $(kl)/(\log N)$ where k is the number of elements per vector (*i.e.*, the dimension), and l is the number of bits per element, usually eight. The preferred system uses 16384-word codebooks, leading to a compression rate for this stage of the pipeline of $(48 \times 8)/(\log 16384) = 384$ bits/14 bits = 27:1. To simplify decoding, the preferred system represents each index using an integral number of bytes, which reduces compression slightly, to 24:1.

2. Entropy Coding

The second stage of the compression pipeline shown in Fig. 11 is an entropy coder designed to decrease the cost of representing high-probability code indices. Since objects are typically rendered or photographed against a constant-color background, the array may contain many tiles that occur with high probability. The preferred system employs gzip, an implementation of Lempel-Ziv coding. See J. Ziv & A. Lempel, "A universal algorithm for sequential data compression," *IEEE Transactions on Information Theory*, IT-23:337-343, 1977. In this algorithm, the input stream is partitioned into nonoverlapping blocks while constructing a

- 25 -

dictionary of blocks seen thus far. Applying gzip to the array of code indices typically provides an additional 5:1 compression. Thus total compression is $24 \times 5 = 120:1$.

3. Decompression

5 Decompression occurs in two stages. In the preferred embodiment, the first stage -- gzip decoding -- is performed as the file is loaded into memory. The output of this stage is a codebook and an array of code indices packed in 16-bit words. Although some efficiency has been lost by this decoding, the light field is still compressed 24:1, and it is now represented in a way that supports random access.

10 The second stage -- dequantization -- proceeds as follows. As the observer moves through the scene, the display engine (see below) requests samples of the light field. Each request consists of a (u, v, s, t) coordinate tuple. For each request, a subscripting calculation is performed to determine which sample tile is being addressed. Each tile corresponds to one quantization vector and is thus represented
15 in the index array by a single entry. Looking this index up in the codebook reveals a vector of sample values. A second subscripting calculation is then performed, providing the offset of the requested sample within the vector. With the aid of precomputed subscripting tables, dequantization can be implemented very efficiently.

20 D. Display

 The final part of the technique according to the invention is a real time viewer that constructs and displays an image from the light slab given the imaging geometry. In this example, the viewer must resample a two-dimensional slice of lines from the four-dimensional light field; each line represents a ray through the
25 eye point and a pixel center as shown in Fig. 12. There are two steps to this

- 26 -

process: 1) computing the (u, v, s, t) line parameters for each image ray; and 2) resampling the radiance at those line parameters.

As mentioned previously, a big advantage of the light slab representation is the efficiency of the inverse calculation of the line parameters. Conceptually the (u, v) and (s, t) parameters are calculated by determining the point of intersection of an image ray with each plane. Thus, any ray tracer can easily be adapted to use light slabs. However, a polygonal rendering system also may be used to view a light slab. The transformation from image coordinates (x, y) to both the (u, v) and the (s, t) coordinates is a projective map. Therefore, computing the line coordinates can be done using texture mapping. The uv quadrilateral is drawn using the current viewing transformation, and during scan conversion the $(u/w, v/w, w)$ coordinates at the corners of the quadrilateral are interpolated. The resulting $u = uw/w$ and $v = vw/w$ coordinates at each pixel represent the ray intersection with the uv quadrilateral. A similar procedure is used to generate the (s, t) coordinates by drawing the st quadrilateral. Thus, the inverse transformation from (x, y) to (u, v, s, t) reduces essentially to two texture coordinate calculations per ray. This approach is cheap and can be done in real time, and is supported in many rendering systems, both hardware and software.

Only lines with (u, v) and (s, t) coordinates inside both quadrilaterals are represented in the light slab. Thus, if the texture coordinates for each plane are computed by drawing each quadrilateral one after the other, then only these pixels that have both valid $u-v$ and $s-t$ coordinates should be looked up in the light slab array. Alternatively, the two quadrilaterals may be simultaneously scan converted in their region of overlap to cut down on unnecessary calculations; for example, this technique may be used in a software implementation.

To draw an image of a collection of light slabs, the two quadrilaterals of each light slab are drawn sequentially. If the set of lines in the collection of light

- 27 -

slabs do not overlap, then each pixel is drawn only once, resulting in increased efficiency. To further increase efficiency, "back-facing" light slabs are culled.

The second step involves resampling the radiance. The ideal resampling process involves first reconstructing the function from the original samples, and then applying a bandpass filter to the reconstructed function to remove high frequencies that may cause aliasing. The resampling process may be approximated by simply interpolating the four-dimensional function from the nearest samples. This is correct only if the new sampling rate is greater than the original sampling rate, which is usually the case when displaying light fields. However, if the image of the light field is very small, prefiltering is applied. This could easily be done with a four-dimensional variation of the standard midmapping algorithm. See L. Williams, "Pyramidal Parametrics," *Computer Graphics (Proc. Siggraph '83)*, Vol. 17, No. 3, July 1983, pp. 1-11.

Figs. 13A to 13C show the effect of nearest neighbor versus linear interpolation on the $u-v$ plane versus quadrilinear interpretation of the full 4D function. Fig. 13A shows the results where no interpolation is used. Fig. 13B shows linear interpolation in $u-v$ only. Fig. 13C shows quadrilinear interpolation in $u-v-s-t$.

Quadrilinear interpolation coupled with the proper prefiltering generates images with few aliasing artifacts. The improvement is particularly dramatic when the object or camera is moving. However, quadrilinear filtering is fairly expensive and dominates the running time of the viewer. In many situations, the sampling rates in the $u-v$ and $s-t$ planes may be different, and then the benefits of filtering one plane may be much greater than the other plane.

- 28 -

E. Experimental Results

Figs. 14A to 14D show images extracted from four light fields. The first is a buddha constructed from rendered images. The model is an irregular polygon mesh constructed from range data. In this example, the input images were generated using RenderMan, which also provided the machinery for computing pixel and aperture antialiasing. The light field configuration was a single slab similar to that shown in Fig. 3A.

A second light field shown in Fig. 14B is a human abdomen constructed from volume renderings. The two tan-colored organs on either side of the spine are the kidneys. In this case, the input images were orthographic views, so a slab with one plane at infinity was employed as shown in Fig. 4C. Because an orthographic image contains rays of constant direction, more input images were generated than in the example shown in Fig. 4A in order to provide the angular range needed for creating perspective views. The images include pixel antialiasing but no aperture antialiasing. However, the dense spacing of input images reduces aperture aliasing artifacts to a minimum.

A third example shown in Fig. 14C is an outward-looking light field depicting a hallway in Berkeley's Soda Hall, rendered using a radiosity program. To allow a full range of observer motion while optimizing sampling uniformity, four slabs were used with one plane at infinity. Thus this example corresponds to a four-slab version of Fig. 4C. The input images were rendered on an SGI RealityEngine, using the accumulation buffer to provide both pixel and aperture antialiasing.

A fourth example shown in Fig. 14D is a light field constructed from digitized images. The scene is of a toy lion, and the light field consists of four slabs as shown in Fig. 3C, allowing the observer to walk completely around the object. The sensor and optical system provide pixel antialiasing, but the aperture

diameter was too small to provide correct aperture antialiasing. As a result, the light field exhibits some aliasing, which appears as double images. These artifacts are worst near the head and tail of the lion because of their greater distance from the axis around which the camera rotated.

Table I summarizes the statistics of each light field shown in Figs. 14A to 14D. Table II provides acquisition parameters for the lion dataset (Fig. 14D). Table III gives the performance of the compression pipeline on two representative datasets. Table IV summarizes the performance of our interactive viewer operating on the lion light field.

In the examples shown, the buddha (Fig. 14A) was compressed using a two-dimensional tiling of the light field, yielding a total compression rate of 45:1. The lion (Fig. 14D) was compressed using a four-dimensional tiling, yielding a higher compression rate of 118:1. During interactive viewing, the compressed buddha is indistinguishable from the original; the compressed lion exhibits some artifacts, but only at high magnifications. Representative images are shown in Figs. 15A and 15B. Based on other experiments at higher rates, we have found that as a general rule, the artifacts become objectionable only above 200:1.

TABLE I: Statistics of the light fields shown in Figs. 14A-14D.

	buddha	kidney	hallway	lion
Number of slabs	1	1	4	4
Images per slab	16x16	64x64	64x32	32x16
Total images	256	4096	8192	2048
Pixels per image	256 ²	128 ²	256 ²	256 ²
Raw size (MB)	256	201	1608	402
Prefiltering	uvst	st only	uvst	st only

In Table II below, distance to object and camera pan per sample are given at the center of the plane of camera motion. Total acquisition time includes longer gantry movements at the end of each row and manual setup time for each of the

- 30 -

four orientations. The aperture diameter is the focal length divided by the F-number.

TABLE II: Acquisition parameters for the lion field.

Camera motion	
translation per slab	100 cm x 50 cm
pan and tilt per slab	90° x 45°
number of slabs	4 slabs 90° apart
total pan and tilt	360° x 45°
Sampling density	
distance to object	50 cm
camera pan per sample	3.6°
camera translation per sample	3.125 c
Aperture	
focal distance of lens	10mm
F-number	f/8
aperture diameter	1.25 mm
Acquisition time	
time per image	3 seconds
total acquisition time	4 hours

Table III describes compression statistics for two light fields. The buddha light field was compressed using two-dimensional tiles of RGB pixels, forming twelve-dimensional vectors. The lion light field was compressed using four-dimensional tiles (two-dimensional tiles of RGB pixels from each of 2x2 adjacent camera positions), forming 48-dimensional vectors. Bytes per codeword index include padding as described above. Peak signal-to-noise ratio (PSNR) is computed as $10\log_{10}(255^2/\text{MSE})$.

TABLE III: Compression statistics for two light fields.

		buddha	lion
	Vector quantization		
	raw size (MB)	50.3	402.7
5	fraction in training set	5%	3%
	samples per tile	2x2x1x1	2x2x2x2
	bytes per sample	3	3
	vector dimension	12	48
	number of codewords	8192	16384
10	codebook size (MB)	0.1	0.8
	bytes per codeword index	2	2
	index array size (MB)	8.4	16.8
	total size (MB)	8.5	17.6
	compression rate	6:1	23:1
15	Entropy coding		
	gzipped codebook (MB)	0.1	0.6
	gzipped index array (MB)	1.0	2.8
	total size (MB)	1.1	3.4
	compression due to gzip	8:1	5:1
20	total compression	45:1	118:1
	Compression performance		
	training time	15 mins	4 hrs
	encoding time	1 mins	8 mins
	original entropy (bits/pixel)	4.2	2.9
25	image quality (PSNR)	36	27

Table IV indicates the performance of the interactive viewer for the lion field. Displayed images are 192 x 192 pixels. Sample extraction includes VQ

decoding and sample interpolation. Display overhead includes reading a mouse, computing observer position, and copying the images to a frame buffer. Timings are for a software-only implementation on a 250 MHz MIPS 4400 processor. As the table shows, interactive playback rates are achieved for reasonable image sizes.

Note that the size of the light field has no effect on playback rate; only the image size matters. Memory size is not an issue because the compressed fields are small.

TABLE IV: Display performance for the lion field.

Display times (ms)	no bilerp	uv lerp	uvst lerp
coordinate calculation	13	13	13
sample extraction	14	59	214
overhead	3	3	3
total	30	75	230

F. Conclusion

In fulfillment of the foregoing objects of the invention, we have now described a new light field representation, the light slab, for storing all the radiance values in free space. Both inserting images into the field and extracting new views from the field involve resampling, a simple and robust procedure. In the above described examples, the sampling density must be high to avoid excessive blurriness. This requires rendering or acquiring a large number of images, which may take a long time and consume a lot of memory. However, denser sample spacing leads to greater inter-sample coherence, so the size of the light field is usually manageable after compression. The resulting system is easily implemented on workstations and personal computers, requiring modest amounts of memory and cycles. Thus, this technique is useful for many applications requiring interaction with three-dimensional scenes.

Many modifications and improvements to the specific illustrative examples

[illegible]

- 34 -

will be apparent to those skilled in the art. For instance, in the above examples, the observer is restricted to regions of space free of occluders. This limitation can be addressed by stitching together multiple light fields based on a partition of the scene geometry into convex regions. If the light fields are augmented to include Z-depth, the regions need not even be convex.

Moreover, in the foregoing embodiments, the illumination must be fixed. However, if interreflections are ignored, this limitation can be addressed by augmenting light fields to include surface normals and optical properties. To handle interreflections, we might try representing illumination as a superposition of basis functions. In this regard, see J. Nimeroff et al., "Efficient Rendering of Naturally Illuminated Scenes," *Proc. Fifth Eurographics Rendering Workshop*, 1994, pp. 359-73. This would correspond in our case to computing a sum of light fields each lit with a different illumination function.

The inventive technique may be compared with depth-based or correspondence-based view interpolation. In these systems, a three-dimensional model is created to improve quality of the interpolation and hence decrease the number of pre-acquired images. With the present approach, a much larger number of images is acquired. At first, this may seem like a disadvantage. However, because of the three-dimensional structure of the light field, simple compression schemes are able to find and exploit this same three-dimensional structure. With our invention, simple four-dimensional block coding leads to compression rates of over 100:1. Given the success of the compression, a high density compressed light field has an advantage over other approaches because the resampling process is simpler, and no explicit three-dimensional structure must be found or stored.

There are many representations for light used in computer graphics and computer vision, for example, images shadow and environment maps, light sources, radiosity and radiance basis functions, and ray tracing procedures. The

- 35 -

present invention suggests that it would be fruitful to study other abstract light representations from first principles. Such reexaminations may in turn lead to further methods for addressing the central problems in these fields, including those identified herein.

5 Another area of future research suggested by this description is the design of instrumentation for acquisition. For example, a large parallel array of cameras connected to a parallel computer could be built to acquire and compress a light field in real time. In the short term, there are many interesting engineering issues in designing and building gantries to move a small number of cameras and lights to
10 sequentially acquire both inward- and outward-looking light fields. This same instrumentation could lead to breakthroughs in both three-dimensional shape acquisition and reflection measurements. In fact, the interaction of light with any object can be represented as a higher-dimensional interaction matrix; acquiring, compressing, and manipulating such representations are a fruitful area for
15 investigation.

 The above-described embodiments of the invention and illustrations thereof are set forth only by way of example. Many modifications and variations beyond those specifically mentioned will be apparent to those skilled in the art without departing from the spirit and scope of the invention.

- 36 -

WHAT IS CLAIMED IS:

1. A method for creating and displaying a scene comprising the steps of:
creating a light field comprising of a collection of images as seen from different
directions and points of view; storing and compressing data corresponding to the
light field; and computing and displaying new images by resampling a subset of the
light field.

2. The method according to claim 1, wherein the light field represents a
virtual scene.

3. The method according to claim 1, wherein the light field represents a real
environment.

4. The method according to claim 1, wherein at least some of the new
images are not contained in acquired data corresponding to the light field.

5. The method according to claim 1, wherein the light field represents an
object, and the step of computing and displaying includes synthesizing new views
looking inward at the object.

6. The method according to claim 1, wherein the light field represents an
environment, and the step of computing and displaying includes synthesizing new
views looking outward into the environment.

7. The method according to claim 1, wherein the step of creating a light
field includes representing a light field in free space by four-dimensional sets of
rays.

- 37 -

8. The method according to claim 7, wherein the four-dimensional sets of rays are represented by respective intersections of the rays with two planes.

9. The method according to claim 8, wherein each of the two planes are placed at an arbitrary position.

5 10. The method according to claim 9, wherein one of the two planes is placed at infinity.

11. The method according to claim 1, wherein the step of creating a light field includes representing a collection of rays by a plurality of sets of parallel planes.

10 12. The method according to claim 8, wherein step of computing and displaying includes generating new views using texture mapping.

13. The method according to claim 1, further comprising the step of decompressing the data prior to displaying.

15 14. The method according to claim 13, wherein the step of storing and compressing includes a two-stage pipeline compression scheme.

15. The method according to claim 14, wherein the two-stage pipeline compression scheme comprises a fixed rate vector quantization step and an entropy coding step.

20 16. The method according to claim 15, wherein the entropy coding step comprises Lempel-Ziv coding.

- 38 -

17. The method according to claim 13, wherein the step of decompressing comprises a decoding stage wherein a file containing compressed data is loaded into memory and a second dequantization stage executed as samples of the light field are requested for display.

5 18. The method according to claim 1, wherein the light field is created from rendered images of a scene.

19. The method according to claim 1, wherein a four-dimensional light field is created from a two-dimensional array of sheared perspective images of the scene.

10 20. The method according to claim 1, wherein a four-dimensional light field is created from a two-dimensional array of fixed field of view perspective images.

21. The method according to claim 1, wherein a four-dimensional light field is created from a two-dimensional array of orthographic images.

15 22. The method according to claim 19, further comprising the step of prefiltering the light field during creation to alleviate aliasing.

23. The method according to claim 20, further comprising the step of prefiltering the light field during creation to alleviate aliasing.

20 24. The method according to claim 21, further comprising the step of prefiltering the light field during creation to alleviate aliasing.

25. The method according to claim 1, further comprising the step of prefiltering the light field during creation to alleviate aliasing, wherein the step of

- 39 -

prefiltering utilizes at least one of the group consisting of a pixel filter and an aperture filter.

26. The method according to claim 1, further comprising the step of prefiltering the light field during creation to alleviate aliasing, wherein the step of prefiltering corresponds to integrating over a pixel of an image of said scene and integrating over an aperture through which the image is received.

27. The method according to claim 1, wherein the light field is created from a digitized image of the scene.

28. The method according to claim 27, wherein the light field is obtained with a computer-controlled motorized gantry with a mounted camera.

29. The method according to claim 28, wherein the gantry is adapted to translate the camera in a plane.

30. The method according to claim 28, wherein an array of images of an object within the scene is recorded from a plurality of equally spaced perspectives.

31. The method according to claim 29, wherein the light field is created by: translating the camera through a grid of positions; panning and tilting the camera at each position at the center of an object within the scene; acquiring images of the object and each position; and reprojecting each image to lie on a common plane.

32. The method according to claim 28, wherein the camera has an aperture which is adjusted to filter an acquired image to prevent aliasing.

33. The method according to claim 29, wherein images collected at different positions are averaged to prevent aliasing.

- 40 -

34. The method according to claim 1, wherein the step of computing and displaying comprises the steps of computing line parameters for each image ray in a slice; and resampling the radiance at those line parameters.

5 35. The method according to claim 34, wherein two-dimensional parameters associated with a plane is computed using a 3x3 linear transformation matrix as used in texture mapping.

36. The method according to claim 34, wherein a nearest neighbor or box filter is used for resampling.

10 37. The method according to claim 34, wherein a linear filter is used for resampling.

38. The method according to claim 36, wherein different filters are used for different dimensions.

39. The method according to claim 37, wherein different filters are used for different dimensions.

15 40. A system for creating and displaying a scene comprising the steps of: means for creating a light field comprising of a collection of images as seen from different directions and points of view; means for storing and compressing data corresponding to the light field; and means for computing and displaying new images by resampling a subset of the light field.

20 41. The system according to claim 40, wherein the light field represents a virtual scene.

- 41 -

42. The system according to claim 40, wherein the light field represents a real environment.

43. The system according to claim 40, wherein at least some of the new images are not contained in acquired data corresponding to the light field.

5 44. The system according to claim 40, wherein the means for creating a light field includes means for representing a light field in free space by four-dimensional sets of rays.

45. The system according to claim 44, wherein the four-dimensional sets of rays are represented by respective intersections of the rays with two planes.

10 46. The system according to claim 45, wherein each of the two planes are placed at an arbitrary position.

47. The system according to claim 46, wherein one of the two planes is placed at infinity.

15 48. The system according to claim 40, wherein the means for creating a light field includes means for representing a collection of rays by a plurality of sets of parallel planes.

49. The system according to claim 45, wherein the means for computing and displaying includes means for generating new views using texture mapping.

20 50. The system according to claim 40, further comprising means for decompressing the data.

- 42 -

51. The system according to claim 50, wherein the means for storing and compressing includes means for implementing a two-stage pipeline compression scheme.

52. The system according to claim 50, wherein the means for decompressing comprises a decoding stage wherein a file containing compressed data is loaded into memory and a second dequantization stage executed as samples of the light field are requested for display.

53. The system according to claim 40, wherein the light field is created from rendered images of a scene.

54. The system according to claim 53, wherein a four-dimensional light field is created from a two-dimensional array of sheared perspective images of the scene.

55. The system according to claim 53, wherein a four-dimensional light field is created from a two-dimensional array of fixed field of view perspective images.

56. The system according to claim 53, wherein a four-dimensional light field is created from a two-dimensional array of orthographic images.

57. The system according to claim 40, further comprising means for prefiltering the light field during creation to alleviate aliasing.




58. The system according to claim 40, wherein the light field is created from a digitized image of the scene.





- 43 -





59. The system according to claim 58, wherein the means for creating a light field includes a computer-controlled motorized gantry with a mounted camera.

60. The system according to claim 40, wherein the means for creating a light field includes a view camera which creates a four-dimensional light field from a two-dimensional array of sheared perspective images of a scene.

5

**FIG._1****FIG._2A****FIG._2B**





4 / 12




FIG. 5




FIG. 7

5 / 12




FIG. 6A

6 / 12



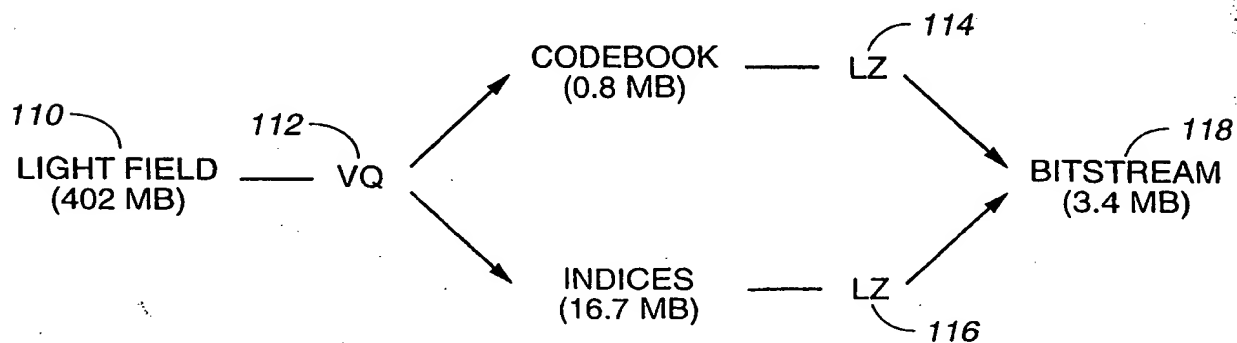


FIG. 6B

7 / 12

**FIG._8****FIG._11**

8 / 12




FIG._9



FIG._13A



FIG._13B



FIG._13C

9 / 12




FIG. 10




FIG. 12

10 / 12






FIG. 14A




FIG. 14B

11 / 12







FIG. 14C

FIG. 14D

12 / 12



FIG. 15B




FIG. 15A